

## Text-to-Scene Conversion

### FIELD OF THE INVENTION

The invention relates to the conversion of text into a three dimensional depiction of the text.

5

### BACKGROUND

At least one Internet service exists where a user can attach his or her text to an image so that the image and text can be e-mailed as an electronic greeting card. The service, provided by BlueMountain.com, relies on a predetermined catalog of images. A user may only add text to the images found in the predetermined catalog. Another system, by Norman Badler at the University of Pennsylvania, uses a text input to drive a closed three-dimensional graphics environment. Badler's system instructs a character to perform certain actions, it will not construct a character or scene based on input text. It is limited to its preexisting environment.

15

### SUMMARY OF THE INVENTION

The invention relates to a method of converting text into three-dimensional scene descriptions, which may then be rendered into three-dimensional images. The invention may generate arbitrary scenes in response to a substantially unlimited range of input text. In one embodiment of the invention, text may be entered into a computer. The text may be passed to a 20 part of speech tagger, which tags the text with grammatical parts of speech. The text may then be parsed. The output of the parser may be converted to a dependency structure. The dependency structure may be semantically interpreted and converted into a scene description.

The scene description may be interpreted into a three-dimensional scene. The three-dimensional scene may be rendered into an image.

## BRIEF DESCRIPTION OF THE DRAWINGS

5 FIG. 1 is a flow diagram indicating the general flow of an embodiment of the invention.

10 FIG. 2 is a flow diagram indicating the general flow of an embodiment of the invention.

15 FIG. 3 illustrates a dependency structure.

20 FIG. 4 is a flow diagram of an embodiment of a depiction module.

25 FIG. 5 is an illustration exemplifying a “canopy area.”

30 FIG. 6 is an illustration exemplifying a “top surface.”

35 FIG. 7 is an illustration exemplifying a “ridge.”

40 FIG. 8 is an illustration exemplifying a “peak.”

45 FIG. 9 is an illustration exemplifying a “base.”

50 FIG. 10 is an illustration exemplifying a “stem.”

55 FIG. 11 is an illustration exemplifying a “cup.”

60 FIG. 12 is an illustration exemplifying a “cap.”

65 FIG. 13 is an illustration exemplifying an “enclosure.”

70 FIG. 14 is an example of a rendering of the text: “*The rose is in the vase.*”

75 FIG. 15 is an illustration exemplifying a “throw ball” pose.

80 FIG. 16 is an example of a rendering of the text: “*John kicks the bowling ball to the cat on the skateboard.*”

85 FIG. 17 is an example of an image of a baby carriage with a tagged handle.

FIG. 18 is an example of a rendering of the text: “*The bird is in the birdcage.*”

FIG. 19 is an example of a rendering of the text: “*The bird is in the birdcage. The birdcage is on the chair.*”

FIG. 20 is an example of a rendering of the text: “*The cowboy is studying mycology.*”

5 FIG. 21 is an example of a rendering of the text: “*The blue daisy is not in the vase.*”

FIG. 22 is an example of a rendering of the text: “*The two cats are on the table.*”

## DETAILED DESCRIPTION OF AN EMBODIMENT OF THE INVENTION

10 A system to generate arbitrary scenes in response to a substantially unlimited range of words is presented herein. The system converts a set of words into three-dimensional scene descriptions, which may then be rendered into three-dimensional graphical images. As used hereinafter, the word “text” describes a set of words that may instantiate an idea or concept, regardless of the form used to express the words (such as, e.g., speech, printed material, visual display, etc.).

15 As an example to illustrate possible inputs and outputs of such a system, consider the sentences “*John said that the cat was on the table. The animal was next to a bowl of apples.*”

An embodiment of the system might construct a picture of a human with a cartoon speech-bubble coming out of his mouth. The speech-bubble might contain an image of a cat on a table.

20 Furthermore, the cat might be next to a bowl containing more than one apple.

In one embodiment, the system may include a set of software programs, mostly written in Common Lisp, but parts of which may be written in C, Java, and Unix shell programs. Other programming languages may be used without departing from the scope of the invention.

FIG. 1 is a flow diagram indicating the general flow of an embodiment of the invention.

Text may be input to the system at block 30. Text may be input by, for example, typing the text into a computer using the computer's keyboard, entering the text via a voice-to-text translation program, entering the text via an optical character recognition system, importing a text file from a memory device, downloading text in any manner to an input port on the system, or entering 5 text in any of the many ways available to those of skill in the art. Text may include one or more words, phrases, or sentences. For ease of illustration, all input text will hereinafter be referred to as "text." At block 32, a linguistic analysis may be performed on the text. Linguistic analysis may include part of speech tagging 31, which tags the text with grammatical parts of speech. Linguistic analysis may also include parsing 33 of the text into a parse tree that represents, for 10 example, the grammatical structure of a sentence. Linguistic analysis may also include converting 35 the parse tree into a dependency structure. A dependency structure typically shows each word in a sentence, and for a given word the dependency structure shows the words that the given word is dependent on and shows the words that depend on the given word. The dependency structure may be semantically interpreted 36 and converted into a high-level scene 15 description.

The high-level scene description may be converted into a low-level scene description in a depiction module 34. The depiction module 34 may convert 39 the high-level scene description to description elements. Description elements may be objects representative of text to be depicted and relationships between those objects. At block 41, depiction rules may be used to 20 assign descriptors to description elements. Several databases may be used to apply depiction rules and assign depictors. Exemplary databases are an environmental database 38, an object database 40, a pose database 42, and an emblem database 44. At block 46, procedures defined by

depictors for transforming the objects representative of the text into a low-level scene description may be executed. At block 48 the system provides an output, which may be, for example, a three-dimensional scene description or a three-dimensional rendered image. The output 48 may include the rendering of a three-dimensional image on a computer screen, the generation of a 5 three-dimensional scene description in a computer file such as a VRML file, the printing of a three-dimensional image on hard copy, or any other three-dimensional image format known to persons of skill in the art.

Environmental database 38 may include environment or settings for images. Such an environment or setting might also be considered as a background upon or in which the text may be portrayed. Object database 40 may include a plurality of three-dimensional models for objects to be included in a low-level scene description. In addition to three-dimensional data, an embodiment may associate additional information with each three-dimensional model, such as a function of an object or its size. Pose database 42 may include poses for actions that may be typically associated with scene descriptions, such as jump, give, and carry. Emblem database 44 may include emblems that might be used to label objects to associate the object with any of a plurality of, for example, occupations or uses. Thus, with an appropriate emblem, a human figure might be either a doctor or a gas station attendant, or any occupation associated with an emblem attached to the human figure; a book might be labeled with an emblem resembling a mushroom to illustrate that the book is used for mycology.

20 FIG. 2 is a flow diagram indicating the general flow of actions typically embodied by the invention. At step 50, text may be entered into a computer. An example of text that might be entered into the system are the sentences "*John said that the cat was on the table. The animal*

was next to a bowl of apples." These exemplary sentences are not meant to be limiting as to the type of text entered. At step **52** the text may be passed to a part of speech tagger, which tags the text with grammatical parts of speech. The text may then be parsed. At step **54**, the output of the parser may be converted to a dependency structure. At step **56** the dependency structure may be 5 semantically interpreted and converted into a scene description. At step **58**, the scene description may be interpreted into a three-dimensional scene. At step **60**, the three-dimensional scene may be rendered into an image.

In one embodiment, steps **50**, **54**, **56**, **58**, and **60** may be implemented in Common Lisp, while steps **58** and **60** may be further implemented on top of, such as, the Mirai® animation system. The embodiment may run on a 450 MHz Pentium® III personal computer under the Windows NT® operating system. A graphics board, such as an Oxygen® VX1 by 3Dlabs® may be used. Step **52** may run on a 194 MHz Silicon Graphics, Inc. model IP25 running the IRIX® 6.5 operating system. Data may be stored locally or remotely. In one embodiment, data may be stored on a Unix® fileserver. Of course, the steps recited above may be implemented using other computer languages/programs without departing from the scope of the invention. Additionally, recitation of computing speeds and types of hardware for computing and storage, as used in the embodiment just disclosed, are not intended as a limitation on the selection of computer hardware or operating systems. Other hardware and operating systems may be used without departing from the scope of the invention.

20 In another embodiment, the system might be run entirely within the same software system. In yet another embodiment, one might eliminate the tagging and parsing phase in favor of a direct construction of the dependency structure.

Steps 52, 54, 56, 58, and 60 are more fully described below.

## Part of Speech Tagging and Parsing

After entry, text may be passed to a part of speech tagger, which tags the text with

5 grammatical parts of speech. In one embodiment, the part of speech tagger used is the MXPOST statistical tagger by Adwait Ratnaparkhi. For the first of the exemplary sentences given above, the result of the part of speech tagging is as given below:

*John/NNP said/VBD that/IN the/DT cat/NN was/VBD on/IN the/DT table/NN*

This result indicates that *John* is a proper noun (NNP), *said* and *was* are past tense verbs (VBD), *the* is an article (DT), *on* is a preposition (IN) (note that in this case *that* is mistagged as a preposition) and *cat* and *table* are nouns (NN). The listing of parts of speech in the above example is illustrative and not meant to be limiting.

The part of speech tagged sentence may be passed to a statistical parser, such as a statistical parser described by Michael Collins, in Head-Driven Statistical Models for Natural Language Parsing (1999) (University of Pennsylvania). The output of a statistical parser is a parse tree that represents the structure of the sentence. The parse tree for the sentence: “*John said that the cat was on the table.*” is represented below:

```
((:HEAD "TOP" "said" 1 1 "node1")
  ((:HEAD "S" "said" 2 2 "node1")
    ((:HEAD "NP" "John" 1 1 "node2")
      ((:HEAD "NPB" "John" 1 1 "node2") ("John" "NNP")))
      ((:HEAD "VP" "said" 2 1 "node1") ("said" "VBD"))
      ((:HEAD "SBAR" "that" 2 1 "node3") ("that" "IN")))
```

```

(:HEAD "S" "was" 2 2 "node4")
(:HEAD "NP" "cat" 1 1 "node5")
  ((:HEAD "NPB" "cat" 2 2 "node5") ("the" "DT") ("cat" "NN")))
(:HEAD "VP" "was" 2 1 "node4") ("was" "VBD")
5  ((:HEAD "PP" "on" 2 1 "node6") ("on" "IN"))
  ((:HEAD "NP" "table" 1 1 "node7")
    ((:HEAD "NPB" "table" 2 2 "node7") ("the" "DT") ("table" "NN")
     ("." "PUNC."))))))))

```

10 The following identifiers have the following meaning: TOP identifies the top of the tree, S stands for embedded sentence, NP stands for noun phrase, NPB identifies a type of noun phrase, VP stands for verb phrase, SBAR stands for complement sentence, PP stands for prepositional phrase, and PUNC. identifies the punctuation mark, in this example the period. This exemplary list is meant to be illustrative and not limiting, the structure of the sentence may be represented in other ways without departing from the scope of the invention.

15 The above exemplary parse tree represents two kinds of information. First, it represents the constituent structure of the sentence, or in other words which words belong together. For example, it represents that *the cat* is a noun phrase (NP), *on the table* is a prepositional phrase (PP), *that the cat was on the table* is an complement sentence (SBAR), and so forth. Second, it represents the headedness of the sentence, or in other words which words are the heads or 20 “central” words of their constituents. For example, it represents that *cat* is the head of its noun phrase (NP) (note at parse tree line 8 the result (:HEAD "NP" "cat" 1 1 "node5")). Similarly, it represents that *was* is the head of the embedded sentence (S) (note at parse tree line 7 the result (:HEAD "S" "was" 2 2 "node4")).

The node numbers listed in the structures are unique identifiers for each node in the parse tree; except that, in the embodiment describer, the node number for a constituent is always the same as the node number of its head. Note, for example, that the verb phrase (VP) “node4” headed by *was* is the head of the embedded sentence (S) that dominates it, which is also labeled 5 “node4”.

### Dependency Structure

In FIG. 2, at step 54, the parse tree may be converted into a dependency structure. A dependency structure typically shows each word in a sentence, and for a given word the dependency structure shows the words that the given word is dependent on and shows the words that depend on the given word. For example, in a phrase, all non-head words and heads of all non-head phrases are dependent on the head of that phrase. FIG. 3 illustrates a dependency structure 62 for the sentence: “*John said that the cat was on the table.*” Arrows 64 point from headwords to words that are dependent on the headwords.

A dependency structure is one possible representation of the semantic relations of a sentence. A dependency structure may enable the focusing of attention on certain dependents, while other dependents may be ignored. For example, one might be interested in interpreting all adjectives that depend upon a noun (*e.g., the large fat green cat*), but ignore the article *the*. In another embodiment, the semantic relations may be computed directly from the parse trees via a tree-transduction rule.

20 The dependency structure may be computed from the parse tree as follows: if a constituent *Y* is the head of a phrase and *X* is a non-head, then in the dependency structure make *head\*(Y)* the head and *head\*(X)* the dependent. *head\** is a recursive function, defined as

follows, where  $head(X)$  returns the head of constituent  $X$ :

1. If  $X$  is a word then  $head^*(X) = X$ ;
2. otherwise  $head^*(X) = head^*(head(X))$ .

Adjustments to the dependency structure derived via this simple algorithm may be

5 computed at this stage. One adjustment may be to share left-dependents between verbs. This is particularly useful for modal verbs such as *may*. The dependency structure for *John may go* would have *John* as a left-dependent of *may*, and *go* as the right-dependent. However, *John* should also be considered a left-dependent of *go*, because that noun phrase is the subject of *go*.

This may be accomplished by allowing *may* to share its left-dependent with the verb that it is the head of, namely *go*. Note that this process of left-dependent sharing may be lexically driven, meaning that it may be performed for those words that specify, in their lexical representation, that they should share their left-dependents. A lexical representation of *may*, which specifies the share-ldependents feature, is shown below:

```
(lexical-entry-macro "may"
  :lemmata '("may")
  :parts '("MD")
  :operations 'share-ldependents
  :semantics (semantics
    :genus "possibility"
    :funcs '(sem-modal-dependent)))
```

20 Also shown in this lexical representation is other relevant information such as the part of speech (MD = "modal"), the lemmata that may be found for this word, and the semantic frame including information about semantic functions invoked by this lexical entry (see section entitled "Semantic Analysis and the Scene Description").

The same dependent sharing is done by the verb *be* so that in the example at hand, *was* shares its left-dependent *cat* with its right-dependent *on*, so that the left-dependent of *on* will now also be *cat*.

Another adjustment may be to combine certain complex words into single words, for 5 easier subsequent manipulation. An example is the complex verb *give up* as in *John gave the money up*. This may be handled in the system by searching the dependency structure for cases where a head word (*give*) is also the first word in a multi-word lexical entry (*give up*), and using a greedy algorithm to capture the dependents (*up*) that match the remainder of that complex entry.

10 In addition to these structural adjustments, lemma information may be added to the dependency representation: a lemma is the “dictionary” form of the word, without inflectional endings. So the lemma of *dog* is itself, *dog*, but the lemma of *dogs* is *dog*; the lemma of *eats* is *eat*; and the lemma of *was* is *be*. This information may be useful because the lexical entries of words may be stored in terms of their lemmata, not in terms of their individual stem forms. In an 15 embodiment of the invention, lemmatization may be accompanied by table lookup: a large list of inflected forms and their associated lemmata may be kept. In another embodiment, one might use a morphological analyzer to compute the lemma of a word. Still other methods of lemmatization are possible without departing from the scope of the invention.

### **Semantic Analysis and the Scene Description**

20 Returning now to FIG. 2, at step 56, the dependency structure may be converted to a scene description. The scene description may be a description of the objects to be depicted in the scene, and the relationships between the objects.

An example of a scene description for the sentence: “*John said that the cat is on the table.*” is given below:

5

10  
15  
20

```
((“node2”
  (:OBJECT (“mr-happy-group”):LEXICAL-SOURCE “John” :SOURCE SELF))

  (“node1”
    (:ACTION “say” :SUBJECT “node2” :DIRECT-OBJECT
      (“node3” “node5” “node4” “node7”):TENSE-ASPECT “PAST” :SOURCE
      “VF:SAY-BELIEVE-THAT-S-FRAME”)
    (“node3” (:STRUCTURAL-NODE “node4” :SOURCE COMPLEMENTIZER))
    (“node5” (:OBJECT (“cat-vp2842”):SOURCE GET-INSTANCES))
    (“node4”
      (:STATIC-RELATION “on” :FIGURE “node5” :GROUND “node7”
        :STRUCTURAL-SOURCE “node6” :SOURCE STATIC-RELATION-PRED))
    (“node7”
      (:OBJECT
        (“table-vp14364” “nightstand-vp21374” “table-vp4098” “pool_table-vp8359”
          “table-vp18802” “chinese_furniture_table-vp23846” “desk_childs-vp41795”
          “desk_w_drawers-vp23679” “card_table_casino-vp42619”
          “ping_pong_table-vp21198” “roulette_table_casino-vp42624”
          “counter_kitchen_island-vp16246” “bathroom_vanity-vp15880”
          “dressing_table-vp6457” “table-vp20056”)
        :SOURCE GET-INSTANCES)))
  ))
```

20

25

In an embodiment, the scene description may be a list of scene description fragments,

each scene description fragment corresponding to a particular node of the dependency structure, and ultimately to the parse tree. Consider “node1”, which is the scene description fragment for the ACTION *say* (note that the designation “node1” corresponds to the node for the structure

headed by *say* in the original parse structure). Node1 is an ACTION, of type “say”, whose SUBJECT is “node2”, DIRECT-OBJECT of the collection of “node3”, “node5”, “node4”, and “node7”, whose TENSE-ASPECT is “PAST”, and whose SOURCE is “VF:SAY-BELIEVE-THAT-S-FRAME”.

5 The SUBJECT, “node2”, corresponds to the noun phrase *John*; it is an OBJECT, whose designation is “Mr. Happy Group”. The OBJECT “node5” consists of a list comprising a single three-dimensional model. Three-dimensional models may be obtained from, for example, Viewpoint Digital, a 3D digital content provider. In an embodiment, the three dimensional model may be a Viewpoint™ model identified as “cat-vp2842”. The designation “vp2842” is a Viewpoint™ catalog number. Similarly, the OBJECT “node7” comprises a list of three-dimensional objects that are consistent with the description *table*. Finally, “node4” is the STATIVE-RELATION “on”, whose FIGURE is the cat, and whose GROUND (the thing the cat is on) is the table. “node3” corresponds to the COMPLEMENTIZER *that*, which requires no further description.

10 The scene description fragments may be derived from the dependency structure by semantic interpretation frames. The appropriate semantic interpretation frames may be found by table lookup, given the lemma of the word in question. These frames may differ depending upon what kind of thing the word in question denotes.

15 For nouns such as *cat* or *table*, an embodiment may use an electronic thesaurus, such as, for example, Wordnet®. The Cognitive Science Laboratory of Princeton University developed Wordnet®. Other electronic thesauruses may alternatively or concurrently be used. The thesaurus provides various kinds of semantic relations between words, the particular information

of interest here being the *hypernym* and *hyponym* relations. Thus, the thesaurus may indicate, for instance, that *cat* is a mammal, or that *table* is a kind of furniture. The three-dimensional model may be keyed, indexed, or otherwise cross-referenced to the electronic thesaurus hierarchy so that a designation such as *cat* will retrieve a list of all cats, the designation *table* will retrieve a 5 list of all tables, the designation *animal* will retrieve a list of all animals, and so forth. In an embodiment, the three-dimensional Viewpoint™ models have been keyed into the Wordnet® hierarchy. Other combinations of three-dimensional objects and electronic thesauruses may be used.

Thus, the semantic interpretation of nouns may be accomplished by finding the node in the electronic thesaurus hierarchy corresponding to the three-dimensional model of the word in question. If there is no such node (*i.e.*, if the word does not correspond to a model in the three-dimensional model database) then a list consisting of just the word may be generated for the object description in the corresponding scene description fragment. For example, if there is no three-dimensional model for the word “bagel,” then a list consisting of just the word “bagel” may be generated for the object description. Eventually, once the scene is rendered, the system may 15 display as three-dimensional extruded text, the word “BAGEL” at the place in the scene where the bagel is to appear.

Noun scene description fragments may also contain information on whether the noun in question is plural, determined from a grammatical part of speech tag produced by the part of 20 speech tagger discussed in the section entitled “Part of Speech Tagging and Parsing.” In the embodiment described herein, “NNS” is the part of speech tag identifying plural nouns. An instance of plural marking appearing in a scene description fragment is shown below:

(“node17”

(:OBJECT (“caramel\_apple-vp22054” “apple-vp8326”) :SOURCE  
GET-INSTANCES :PLURAL T))

For personal names such as *John*, a list of male and female first names may be mapped,

5 as appropriate, to one of a male or female three-dimensional humanoid figures. In an embodiment, two three-dimensional humanoid figures may be identified. “Mr. Happy” may be used to depict males, while “Mrs. Happy” may be used to depict females. Of course, other characters may also be available. Additionally, images may be mapped or added to the three-dimensional characters to depict any number of likenesses (*e.g.*, celebrities, political figures, friends, pets).

10 Spatial prepositions such as *on* may be handled by semantic functions that look at the left- and right-dependents of the preposition, and construct a scene description fragment from these. The semantic entry for *on* is given below:

15 (semantics  
:genus “on”  
:funcs '(SEM-TIME-RELATION-ATTR  
SEM-STATIVE-RELATION-PRED  
SEM-STATIVE-RELATION-ATTR  
SEM-INTRANSITIVE-STATIVE-RELATION)

20 The semantic functions (:funcs) may be checked in sequence, with the first one that matches being allowed to construct the scene description fragment. SEM-TIME-RELATION-ATTR will match if the right-dependent of *on* is a time expression (*e.g.*, *Tuesday*). SEM-STATIVE-RELATION-PRED will match cases where there is both a left-dependent and a right-dependent. SEM-STATIVE-RELATION-ATTR will match cases where there is no left-

dependent but there is a head, for example, such a case is *the cat on the mat*, where *cat* is the head and *on* is the dependent. Furthermore, SEM-INTRANSITIVE-STATIVE-RELATION handles the case of intransitives, such as *on* as in *John ran on*. In an embodiment, these semantic functions are implemented as Common Lisp methods. However, other embodiments may

5 implement the semantic functions as functions in some other programming language, or in any other computational method known to those of skill in the art.

Most verbs may be handled not by semantic functions, but rather may be handled by semantic frames. The semantic entry for *say* is shown below:

10 (SEMANTICS :TYPE NIL :GENUS say :SPECIES NIL  
 :FUNCS NIL  
 :VERB-FRAMES  
 ((VERB-FRAME  
 :NAME SAY-BELIEVE-THAT-S-FRAME  
 :REQUIRED (SUBJECT THAT-S-OBJECT)  
 :OPTIONAL (ACTIONLOCATION ACTIONTIME))  
 (VERB-FRAME  
 :NAME SAY-BELIEVE-S-FRAME  
 :REQUIRED (SUBJECT S-OBJECT)  
 :OPTIONAL (ACTIONLOCATION ACTIONTIME))  
 20 (VERB-FRAME  
 :NAME GIVE-FRAME-NP-TOOBJ  
 :REQUIRED  
 (SUBJECT-VOLIT  
 DIRECT-OBJECT  
 INDIRECT-OBJECT)  
 25 :OPTIONAL

(ACTIONLOCATION INSTRUMENT ACTIONTIME))

### (VERB-FRAME

:NAME NP-V-NP

### :REQUIRED (SUBJECT DIRECT-OBJECT)

:OPTIONAL

### (ACTIONLOCATION

## INSTRUMENT

ACTIONTIME))))

5

10 This semantic entry includes a set of verb frames, each of which defines the argument  
structure of one “sense” of the verb *say*. For example, the first verb frame identified as SAY-  
BELIEVE-THE-S-FRAME, has as a required argument a SUBJECT and a THAT-S-OBJECT,  
or in other words an expression such as *that the cat is on the table*. Optional arguments include  
an action location (e.g., **John said in the bathroom that the cat was on the table**) and an action  
time (e.g., **John said yesterday that the cat was on the table**). Each of these argument  
specifications may cause a function to be invoked. The function may check the dependencies of  
the verb for a dependency with a given property, and assign such a dependency to a particular  
slot in the scene description fragment. For example, the THAT-S-OBJECT argument invokes a  
function that finds a dependent whose lemma is *that* and whose grammatical category is SBAR,  
20 and assigns it to the direct object slot of a scene description fragment that corresponds to a verb.

Note that in the scene description below:

```
(“node1”
  (:ACTION “say” :SUBJECT “node2” :DIRECT-OBJECT
  (“node3” “node5” “node4” “node7”) :TENSE-ASPECT “PAST” :SOURCE
  “VF:SAY-BELIEVE-THAT-S-FRAME”))
```

the DIRECT-OBJECT actually consists of several nodes: “node3”, “node5”, “node4”, and “node7”. This is an example of a list of nodes in the scene description corresponding to all elements of the embedded sentence *that the cat is on the table*. This may be provided as a convenience for the depiction algorithms described in the section entitled “Interpretation of the

5 Scene Description into a Three-Dimensional Image.”

The next frame SAY-BELIEVE-S-FRAME handles sentences such as, for example, John *said the cat is on the table* (i.e., without *that*). The frame GIVE-FRAME-NP-TOOBJ handles sentences such as, for example, John *said “bonjour” to Mary*. The frame NP-V-NP handles sentences such as, for example, *John said “bonjour”*.

In addition to semantically interpreting words that denote particular objects, actions, or relations, an embodiment may also interpret anaphoric or coreferring expressions. Simple pronominals like *he* or *she*, may be interpreted by searching through the context to find an appropriate coreferent (where appropriate may include matching on, for example, number and gender features).

15 Nouns can also corefer, as in the following example:

*John said that the cat was on the table. The animal was next to a bowl of apples.*

While it is not strictly required that *the animal* denote *the cat* mentioned in the first sentence of the above example, the coherence of the discourse depends upon the reader or listener making that connection. In an embodiment, such associations may be handled by noting that in the

20 thesaurus’s hierarchy, the denotations of *cat* are a subset of the denotations of *animal*, and computing that the noun phrase might corefer with the previously mentioned *cat*. In the case of multiple sentences, a list of scene descriptions may be returned and passed to a depiction module

as described below in the section entitled “Interpretation of the Scene Description into a Three-Dimensional Image.”

Consider again the following example sentences:

“*John said that the cat was on the table. The animal was next to a bowl of apples.*”

5 In an embodiment, two scene descriptions may be generated; a first scene description for the first sentence and a second scene description for the second sentence. Each scene description, as discussed earlier, may include multiple scene description fragments. In the second scene description, the object corresponding to the *animal*, may be described by a list of possible animals from a set of three-dimensional models; the object also may contain an annotation for a POSSIBLE-COREFERENT. In this example, the POSSIBLE-COREFERENT would be the node corresponding to the *cat*. Thus, the description module (see section entitled Interpretation of the Scene Description into a Three-Dimensional Image) may make use of POSSIBLE-COREFERENT information. Instead of adding a new animal and putting it next to a bowl of apples, the description module may put the already existent cat next to a bowl of apples.

10 The following is one possible example of a scene description for the sentence: “*The animal was next to a bowl of apples.*”

15 ((“node12”  
 (:OBJECT  
 (“pig\_baby-vp1356” “calf-vp1322” “owl-vp11719” “eagle-vp1334”  
 “bird\_standing-vp1316” “atlantic\_puffin-vp13720” “albatross-vp11963”  
 “penguin-vp1354” “penguin-vp7324” “duck-vp1331” “black-swan-vp11716”  
 “ostrich-vp11975” “bird-vp21815” “toucan-vp11967” “chicken-vp3952”  
 “sparrow-vp21210” “manta\_ray-vp13438” “shark-vp24472” “mandarin\_fish-vp14306”  
 “sea-horse-vp23288” “angelfish-vp13793” “marlin-vp1353” “goldfish-vp3982”

“piranha-vp6615” “salmon-vp23287” “mandrill\_male-vp13724” “baboon-vp7801” “chimp-vp2577” “cat-vp2842” “cheetah-vp7811” “tiger-vp1364”  
 “lion\_standing-vp1350” “panda-vp11862” “dog\_standing-vp5041” “dog\_sitting-vp5092”  
 “dog-vp23283” “doberman-vp5851” “republican\_elephant-vp21783” “elephant-vp1338”  
 5 “walrus-vp13713” “dolphin-vp6421” “rabbit-vp23286” “bunny-vp1320”  
 “beaver-vp13721” “rat-vp6423” “chipmunk-vp7813” “rhinoceros-vp7819”  
 “zebra-vp11968” “horse-vp8105” “horse-vp1347” “donkey-vp24215” “camel-vp4014”  
 “giraffe-vp11957” “deer\_mule-vp3614” “caribou-vp13722” “goat-vp2183”  
 “bull-vp13525” “cow\_jointed-vp2234” “bison-vp7807” “night\_glider-vp13892”  
 10 “kangaroo-vp40097” “platypus-vp13726” “frog-vp4411” “fetus\_20-week-vp2870”  
 “pterodactyl-vp1811” “diplodocus-vp1805” “stegosaurus-vp1812” “triceratops-vp1818”  
 “allosaurus\_jointed-vp5489” “brontosaurus-vp8099” “snake-vp24586” “turtle-vp40115”  
 “turtle-vp2230” “starfish-vp4144” “earthworm-vp11587” “snail-vp13117”  
 “octopus-vp11731” “shell\_nautilus-vp40105” “squid-vp23290” “dragonfly-vp21170”  
 15 “grasshopper\_south\_african-vp12855” “beetle\_hercules-vp12649”  
 “beetle\_silver-vp24543” “ant-vp5536” “butterfly-vp14800” “fly-vp13014”  
 “mosquito-vp4767” “crab-vp12487” “crawfish-vp11867” “lobster-vp21752”  
 “spider-vp21211” “scorpion\_imperial-vp12520”)  
 :SOURCE GET-INSTANCES :POSSIBLE-COREFERENT “node5”))  
 20 (“node11”  
 (:STATIVE-RELATION “next to” :FIGURE “node12” :GROUND “node15”  
 :STRUCTURAL-SOURCE “node13” :SOURCE STATIVE-RELATION-PRED))  
 (“node15”  
 (:OBJECT (“bowl-vp41903” “fruit\_bowl-vp24864”):SOURCE GET-INSTANCES))  
 25 (“node16”  
 (:STATIVE-RELATION “of” :FIGURE “node15” :GROUND “node17”  
 :STRUCTURAL-SOURCE “node16” :SOURCE STATIVE-RELATION-ATTR))  
 (“node17”  
 (:OBJECT (“caramel\_apple-vp22054” “apple-vp8326”):SOURCE

GET-INSTANCES :PLURAL T)))

The scene description above illustrates another property of scene descriptions, namely that they may be only partly specified as to the precise semantic relation expressed by the elements. Note “node16”, for the stative relation *of*, which relates *bowl* and *apples*. This gives 5 essentially no information about the relation between *bowl* and *apples*, other than the fact that they are related. The depiction module handles the relation -- in this case that the apples are in the bowl rather than, say, on top of the bowl.

### Interpretation of the Scene Description as a Three-Dimensional Scene

#### Depiction Module Overview

Returning to FIG. 2, at step 58, the high-level scene description, generally as described above, may be passed to a Depiction Module for interpretation of the high-level scene description as a three-dimensional scene. The output of the Depiction Module may be a low-level three-dimensional scene description, such as a low-level scene description in Virtual Reality Modeling Language (VRML) format, or a low-level scene description in any format acceptable to any number of three-dimensional rendering packages or video games, or a rendered three-dimensional image, generally representative of the originally entered text.

FIG. 4 is a flow diagram of the operations that may occur within the Depiction Module. At step 66, description elements may be made, resolved, and sorted based on depiction type and 20 object dependencies. Description elements may be thought of as represent the meaning of the originally entered text; description elements generally correspond to the nodes in the high-level scene description generated by the linguistic module. At step 68, depictors may be assigned to

each description element using depiction rules. At step 70, referenced three-dimensional objects may be read into memory. At step 72, each assigned depictor may be run to incrementally construct a scene. At step 74, supporting objects, ground planes, and lights may be added. At step 76, a user may adjust the viewpoint from which the user views the final rendered three-dimensional image. At step 78, the image may be rendered. It is understood that steps may be combined or the order of steps may be altered without departing from the scope of the invention.

5 Each of the above-identified steps are more fully described below.

### Description Elements, Depictors, and Depiction Rules

10 The linguistic analysis may provide a high-level scene description (hereinafter a “scene description”). A scene description may include a list of nodes (see, for example, the scene description for the sentence: “*The animal was next to a bowl of apples.*”). Each node may be considered a description element. Each description element has a type, which may represent various component meanings inherent in the input text. Examples of description element types may be, for example, OBJECT, ACTION, ATTRIBUTE, STATIVE-RELATION, TIMESPEC, CONJUNCTION, POSSESSIVE, NEGATION, and CARDINALITY. Other description element types may also be available. Additionally each different type of description element may have an object to which it refers. For example, in an embodiment, a STATIVE-RELATION has a FIGURE and a GROUND. In the previous example using the sentence: “*The animal was*

15 *next to a bowl of apples.*”, a STATIVE-RELATION is node16, which has a FIGURE, which is node15, which is the OBJECT *bowl*, and a GROUND, which is node17, which is the OBJECT *apple*.

20

Depictors may be sets of objects and sets of parameters, which may be applied to them.

That is, depictors may be considered as parameterized objects. The depictors may include the reference of the objects to be modified, the parameters to be used in modifying the objects, and the set of procedures for the modification. A depictor may, for example, set the color of another object as in, for example, an ATTRIBUTE description element. Such a depictor may work by calling the software module that adds color to an object. A depictor may also modify the shape of an object, as for example, if the ATTRIBUTE “mangled” were to be used. As another example, a depictor may also make a character swing a baseball bat at a ball as in, for example, an ACTION description element.

Depiction rules may be assigned to description elements and may be used to generate a set of depictors for their assigned description element. The depictors may then be applied in sequence to actually create the scene. In other words, the set of constraints imposed by the depictors may be solved to achieve a satisfactory low-level scene description. The set of depiction rules may be extensible.

Depiction rules may be assigned by first matching the description element type (e.g., ACTION, ATTRIBUTE, OBJECT) and then a particular instance of the scene description. For example, instances of ACTION may be “run” or “throw,” an instance of ATTRIBUTE may be color, and an instance of OBJECT might be to portray the noun *government* as an image of the U.S. Capitol Building. There can be many depiction rules for a given description element and instance. After matching a set of depiction rules, each candidate depiction rule may have a filter applied to it. This filter may be identified by the word “TEST”. TEST may be a piece of code that tests conditions from the entirety of the description. TEST may compute, for example,

whether a certain action is appropriate for an object.

In the first example below, the depiction rule will be considered if either of the actions “kick” or “punt” are depicted. Furthermore, this particular depiction rule is an example of a depiction rule that might be used when there is no path or specified trajectory. An example of a 5 sentence that indicates no path or specified trajectory might be *John kicked the ball*, as opposed to *John kicked the ball over the fence*. This exemplary depiction rule also checks to see that there is a direct object (in this case “*ball*”) and that the size of the direct object is larger than four feet. If the object is smaller than four feet, then a second, possibly less restrictive, depiction rule may be used. Of course, the parameters evaluated by the depiction rule may be changed without departing from the scope of the invention.

10  
15  
20

```
(define-depiction (:action ("kick" "punt") "in place")
  :test (and (not opath)
    direct-object
    (> (find-size-of direct-object) 4.0))
  :fobjects (list (make-pose-depictor "kick" :actor subject)
    (make-spatial-relation-depictor "behind" subject direct-object)))

(define-depiction (:action ("kick" "punt"))
  :fobjects (make-path-verb-depictor subject t 3.0 "kick ball" direct-object opath
  20 :airborne-figure-p t))
```

The function MAKE-PATH-VERB-DEPICTOR, used in the second depiction rule above, creates a set of depictors that depict objects on paths. MAKE-PATH-VERB-DEPICTOR depicts the subject of the sentence “subject”, determines if the subject is to be displayed as being upright “t”, places the subject three feet behind the object (*i.e.*, the ball) “3.0”, uses the “kick ball” pose,

and makes the object airborne. Using depictors, any number of complex objects can be depicted.

Attributes may also have depiction rules. Examples of such depiction rules follow. This list is meant to be illustrative and not restrictive. Other depiction rules are possible.

If an object is described as, for example, “huge”, the object may be scaled from its default

5 size by a factor of four. Other scaling factors are, of course, available. If an object is described as, for example, “spinning,” a spinning icon object (*e.g.*, an arrow curling in a circle) may be situated above the object to suggest spinning motion. It may also be possible to depict motion by causing the object to rotate about an axis. Note that in the example code below, “descriptor” indicates a “description element.”

10  
11  
12  
13  
14  
15

```
(define-depiction (:attribute "huge")
  :fobjects (list object)
  :fdepict (3d:scale-local object 4.0))
```

20

```
(define-depiction (:attribute "spinning")
  :descriptor-role :relation
  :fobjects (list (find-object "spinning icon") object)
  :fdepict (let ((figure (first (get-objects descriptor)))
    (ground (second (get-objects descriptor))))
    (stack-up figure ground :use-surfaces-p nil
      :extra-space
      (find-height-of ground)))
    ;; if the object is repositioned, the spinning icon must move with it
    (attach-for-depiction figure ground)))
```

25

In the example below, TEST computes if the specified attribute is a color. If so, the

function COLORIZE-OBJECT is called to change its the object's color. FIND-COLOR-VALUES identifies if the color is defined in a database of colors.

```
(define-depiction (:attribute nil)
  :test (find-color-values descriptor-name)
  5 :fobjects (list object)
  :fdepict (loop for i in (listify object)
  do (colorize-object i descriptor-name)))
```

In the example below, the test computes if the specified attribute is a facial expression

10 (e.g., "smile"). If so, it calls the function GOTO-FACIAL-EXPRESSION on the object (usually  
a human head) to change its shape appropriately.

```
(define-depiction (:attribute nil "facial expressions")
  :fobjects (list object)
  :test (when (and object descriptor-name)
  15 (find-facial-expression object descriptor-name))
  :fdepict (goto-facial-expression object descriptor-name))
```

## Generating Depictors

Depiction rules may be used to generate depictors for description elements. Depictors  
20 may be responsible for the instantiation of a description element. For example, a depitor for the  
action "eat" might, when applied, place a three-dimensional character in a pose to appear as if the  
three-dimensional character was holding a spoon to its mouth. Note also that more than one  
depitor may be assigned to a single description element. For example, the action of "*John*  
*throws the ball to Mary.*" might generate depictors that could generically be expressed as

follows:

1. John is in an overhand-throwing pose, holding a ball.
2. John is facing Mary
3. Mary is facing John - TENTATIVE
- 5 4. John is 10 feet behind Mary - TENTATIVE

### Transduction Rules

These may be used to modify a set of description elements and/or depictors based on some criteria. In an embodiment, after having assigned depictors to each description element, transduction rules may be applied to find conflicts and add implicit constraints.

### Conflicts

In the example above, depictors 3 and 4 are labeled as "TENTATIVE". This indicates that the depictor in question might be used to set a default; other depictors (from other depiction elements) may override that default. For example, if the input text was "John throws the ball to Mary. John is on the porch. Mary is on the sidewalk." then depictor 4 would be overridden by depictors that put John on the porch and Mary on the sidewalk.

### Constraint Satisfaction

20 Objects may sometimes be multiply constrained. For example, consider "*The plate is to the left of the fork. The plate is to the right of the napkin.*" If the plate is moved in the sequence presented by the two preceding sentences, then the plate will first be put to the left of the fork

and then to the right of the napkin. Additionally, if there is no specification as to where the fork or napkin is initially placed, an unpredictable result is likely. Multiple constraints can be resolved by maintaining previously enforced constraints. In other words, the first constraint may be identified and maintained when the second constraint is implemented. So, in one

5 embodiment, the constraint that the plate is to the left of the fork is maintained when the plate is subsequently moved to the right of the napkin, thus causing the fork to move with it. Other, more sophisticated, constraint satisfaction methods can also be used.

**Implicit constraints.** Implicit constraints are those constraints which may be imposed on objects because of the objects' usage in the context of the entered text. Consider the sentences: "The lamp is on the table. The glass is next to the lamp." It may be preferable not to have the glass floating in the air next to the lamp. Instead, the glass would preferably be put on the table. Therefore, an implicit constraint may be implemented which provides that "If X is next to Y, and X is not already on a surface, and X is not an airborne object (e.g., a helium balloon), then place X on the same surface as Y." Other implicit constraints may additionally be implemented.

10 FIG. 18 is an illustration 132 of one possible rendering of the sentence "*The bird is in the birdcage.*" In this scene, the implicit constraint concerning the objects *bird* and *birdcage* might be that the bird is placed so that its feet are on the floor of the birdcage. Therefore, in this example, the "enclosure" spatial tag for the birdcage and the "base" spatial tag for the bird would be identified and the bird's base would be placed on the floor of the birdcage's enclosure.

15 FIG. 19 is an illustration 134 of one possible rendering of the sentences "*The bird is in the birdcage. The birdcage is on the chair.*" This rendering illustrates the implicit constraint as applied to the positioning of the bird within the birdcage and as applied to the positioning of the

birdcage on the chair.

## Object Database

**3D Models.** A three-dimensional object library may, for example, include models for air

5 transport, anatomy, animals, architecture, characters, geography, ground transport, household, industrial, military, miscellaneous, occupational, space, sports and hobbies, vegetation, and water transport. This listing is meant to be illustrative and not limiting. In one embodiment the

Viewpoint Digital™ “Gold” library of pre-defined three-dimensional polygonal models are used.

An online listing of these models can be found at [www.viewpoint.com](http://www.viewpoint.com). Types of models other than polygonal models may be used; for example non-uniform rational B splines (NURBS) may be used without departing from the scope of the invention. Furthermore, it may be possible to use three-dimensional models from other suppliers and to combine models from multiple suppliers. Additionally, it may not be necessary to use every object modeled in a given library.

For example, in one embodiment approximately 1,200 three-dimensional models, out of approximately 12,640 are utilized.

The three-dimensional object database may be extensible, therefore, in addition to pre-defined three-dimensional models, users can add their own models to the database. A user may model his or her own models using a program such as Mirai™. In addition to the three-dimensional data, an embodiment may associate additional information with each three-dimensional model.

**Parts.** These may be named collections of faces, or surfaces, on the models that represent certain regions. For example, the headlights, roof, and windshield of a car might be

different parts.

**Color Parts.** These may be sets of parts to be colored. Color parts may be used when the model is specified by the text as having a particular color. Therefore, for example, for *the blue flower*, the petals of the flower could be colored blue, not the stem. If no color parts are 5 specified, then the largest part may be colored.

**Opacity Parts.** These may be parts which get a default transparency (e.g., the glass part of a framed window).

**Default size.** All objects may be given a default size. The default size may be expressed in a common unit of measure, for example feet.

**Spatial Tags.** Spatial tags may be areas on and around objects used to spatially arrange and juxtapose objects together. The spatial tag may be marked with a space-filling three-dimensional object, representing the borders of the spatial tag. Spatial tags may be used when depicting, for example, the “in” or “on” spatial relationship. FIGS. 5 through 13 illustrate some spatial tags.

10 FIG. 5 is an illustration exemplifying a “canopy area” 80, which in this illustration is the area under the substantially horizontal portion of a lounge-type chair 82. FIG. 6 is an illustration exemplifying a “top surface” 84, which in this illustration is the area substantially directly above the seat of a chair 86. FIG. 7 is an illustration exemplifying a “ridge” 88, which in this illustration is represented as the lengthwise edge substantially directly above the crest of the roof 15 20 90 covering a well 92. FIG. 8 is an illustration exemplifying a “peak” 94, which in this illustration is a point having a substantially zero cross-sectional area and lying substantially directly above the spout 96 of a water bottle 98. FIG. 9 is an illustration exemplifying a “base”

100, which in this illustration is the area at substantially the lower-most point of a weather vane  
102. FIG. 10 is an illustration exemplifying a “stem” 104, which in this illustration is the area  
typically associated with the stem of a flower 106. Other examples of a stem include, for  
example, a stick on a lollipop, or a vertical support for a stop sign. FIG. 11 is an illustration  
5 exemplifying a “cup” 108, which in this illustration is the area substantially enclosed by the  
outline of a canoe 110. FIG. 12 is an illustration exemplifying a “cap” 112, which in this  
illustration is the area substantially below and surrounded by the portion of a pirate’s hat 114  
wherein the wearer’s head would be located. FIG. 13 is an illustration exemplifying an  
“enclosure” 116, which in this illustration is the area substantially bounded by the interior of a  
10 birdcage 118. Of course, it will be recognized that spatial tags may be associated with objects  
other than those exemplified in the above listed illustrations. For example, the interior area of  
the ashtray functions as a “cup” to contain whatever is placed into the ashtray. Additionally, the  
above list is not meant to limit the number or type of spatial tags. There may be other types of  
spatial tags, for example a spatial tag indicating a “wall” may provide a surface upon which to  
15 push an object. For example, a car may have a wall along the outside surface of its trunk to  
enable a character to push the car by pushing on the wall of the trunk. A wall spatial tag may  
also be used to position an object on a vertical surface, such as when a character hangs a picture  
on a wall.

**Functional Properties.** These may be properties used by the depiction rules to  
20 determine if an object can behave in a given way. For example, all objects that function as land  
vehicles might be marked as such. Then the depiction rule for the verb “ride” may select among  
these to choose a vehicle, as for example, in the sentence *John rides to the store*. The list of

functional properties may be open. The list of functional properties may be stored in one or more databases. The list below is illustrative of the features being marked in one embodiment.

However, the list below is only a small fraction of those features that may be used by the depiction rules:

5 :HARD :SOUR :TRAIN :SALTY :SEPARATOR :COLD-WEATHER-  
THING :CONDUIT :TROPICAL :MONEY :FLATTENER :SIGNAL  
:HOUSEHOLD-ITEM :SUBMARINE :GENERATOR :POKER :PATH  
:CONVEYER :TERRAIN :TRAP :CUTTER :SOFT :TABLE :BEAUTIFUL  
:SEAT :PUSHER :CONTROLLER :INSECT :MIXER :PORTAL  
10 :MOUTHWARE :RECORDER :SUCTION :LIFTER :PUSH-VEHICLE  
:LIGHT-SOURCE :FOOT :ALARM :ROLLER :POINTY :SIGN  
:DINOSAUR :TRUCK :HOUSEWARE :TARGET :HEAT-SOURCE  
:COLD :SWITCH :DEFUNCT :TRIGGER :HAT :PLANT :BED  
:PROJECTILE :FURNITURE :CELEBRITY :SCARY :FRUIT :TREE  
15 :LOCATION :SLIPPERY :HUMAN :MARKER :DISPLAY :RECEPTOR  
:METER :WET :SUPPORT :WEAPON :GUN :MUSICAL-INSTRUMENT  
:DESTROYER :SOUND-SOURCE :VEHICLE :BUILDING :MONUMENT  
:SWEET :GRIPPER :FISH :HOT :SURFACE :HITTER :SOURCE  
:DECORATION :UGLY :BIRD :INFORMATION :TABLEWARE  
20 :COLORFUL :TOY :CELEBRATORY :GROUNDPLANE :TERRAIN  
:CONNECTOR :DANGEROUS :GENERIC :CAR :HEAVY :AIR-  
VEHICLE :GUARDIAN :UNKNOWN :EXOTIC :CATCHER  
:VEGETABLE :CONTAINER :BASIN :SINK :CHAIR :VIEWER :BOAT  
:FLOWER :GAMEWARE :FOOTWARE :DEVICE :SHIELD :PART  
25 :FOOD :ANIMAL :TOOL :CLEANER :BODYWARE

**Rendering Attributes.** Other visual properties of objects can be represented, like

shininess, overall transparency, etc.

**Shape displacements.** Some objects, like human faces, can change shape (e.g., smiling, eyes closed, frowning, etc.). These shape displacements may be associated with the object and may be used to depict emotions or other states of the object.

5

## Spatial Relations

Spatial relations may be depicted using the spatial tags, as discussed above, of associated objects. Some examples:

For *The bird is on the cat*, we find a “top surface” tag for the cat and a “base” tag for the bird. Preferably, the cat’s top surface is on its back and the bird’s base is under its feet. Identification of the spatial tags allow the bird’s feet (base) to be positioned on the cat’s back (top surface). If, for example, the base is too big to fit on the top surface, it could be scaled.

For *The cat is under the chair*, a “canopy” tag for the chair is identified and the cat may be positioned substantially within the boundary of the canopy spatial tag. If, for example, the cat is too large to fit, then either the cat can be scaled or the chair can be put on the back of the cat.

FIG. 14 is an illustration 120 of one possible rendering of the sentence “*The rose is in the vase.*” In this rendering, the “cup” tag for the vase and the “stem” tag for the rose would be identified and the rose’s stem would be placed within the vase’s cupped opening.

These examples are not meant to be an exhaustive list, but rather to illustrate the manner 20 in which spatial tags may be used to depict spatial relations.

## Poses and Templates

In three-dimensional (3D) graphics, human characters are usually represented by a 3D mesh (skin) overlaying a bone control structure. Each vertex in the skin is typically linked to a bone, so that when that bone moves, that part of the skin will move correspondingly. The bones 5 are typically configured in a tree structure, starting with a “root,” usually between the hips. Each bone typically has a length and a rotation angle, representing the bone’s orientation with respect to a parent bone. The total set of bone lengths and rotations is typically called a pose. When the bones are configured into a given pose, the skin may be automatically moved into the corresponding configuration.

10 An extensible library of poses for actions such as, for example, “jump,” “give,” and “carry” may be used in an embodiment of the invention. In one embodiment, over two hundred such poses are used; the number may grow to the thousands and beyond. The notion of poses may be extended to not only consist of poses associated with a humanoid skeleton but also poses associated with the relative size and placement of a set of associated objects. Associated objects 15 may be classified as, for example, “instruments,” “targets,” and “fixtures,” however, other associated objects are also possible.

15 A pose template can be thought of as a representation of a specific form of an action or other description. Therefore, for example, the pose template “swing bat at ball” may represent a humanoid form swinging a representation of a baseball bat at a spherical object representing a 20 baseball. The baseball is the target; the bat is the instrument. These associated objects (e.g., target, instrument, etc.) may be used to make object substitutions in the depiction module. The depiction module will be discussed below. Pose templates can also involve more than one

humanoid form, as in for example, a swordfight. FIG. 15 is an illustration 122 of the “throw ball” pose template.

**Pose template substitution.** The associated objects in pose templates may be represented by generic objects and sets of constraints. Therefore, for example, in the “throw ball” 5 pose template of FIG. 15, a generic sphere 124 represents the ball. If the input sentence is “*John threw the watermelon*”, a watermelon would be substituted for the generic sphere 124. The watermelon can either be substituted as-is, or scaled to approximately the same size as the generic sphere 124. Each pose template may specify whether to scale substituted objects, and if so, what constraints should be met. For example, the position of the substituted object can be manipulated so that the bottom of each object rests on the same plane. FIG. 16 is one possible rendering 126 of a sentence using a kick pose: “*John kicks the bowling ball to the cat on the skateboard.*”

**Partial poses.** Some poses involve the entire object. Others may require only a portion of an object. Therefore, partial poses may be used when one part of an object acts independently of another part of the object. Partial poses may also be combined. For example, a humanoid form running primarily uses the lower body in its pose. In addition, throwing may only require the upper body to be in a pose. Therefore, throwing a ball while running may be depicted by combining the two partial poses. Partial poses need not be restricted to upper and lower halves of an object. A partial pose may be used, for example, to present various poses for a hand on an 20 arm, where the hand acts independently of the position of the arm.

**Inverse Kinematics.** Inverse kinematics (IK) is a 3D graphics technique for specifying the position of an end effector (joint) of a skeleton and having other bones automatically adjust

their angles to reach that position. The typical example is moving the wrist to a target position with IK which causes the elbow and shoulder joints to rotate in the right ways so the wrist reaches the destination. This technique may be used in an embodiment to make characters point and look at objects. IK may be more flexible than poses because the target objects do not have to

5 be in a predetermined place. The skeleton will point at any target.

IK can also be layered on top of existing poses to make the character's body position better correspond with other objects in the environment. An example is with the pose for "push."

Consider pushing a lawnmower, a car, or a baby carriage. Because the various objects have handles and surfaces at different heights, no single body pose can cover them all. (The hands need to touch the correct part of the object.) To take care of this, the character may first be put behind the object in a generic push pose appropriate for that type and size object. Then the hands may be moved using inverse kinematics to the handle or surface of the object. Note that this technique assumes the "push surface" or "handle" of the object is tagged so that the target position for the IK can be specified. FIG. 17 is an image of baby carriage 128 with a tagged handle 130.

## Environment/Setting

The environment or setting of the scene may be specified by the entered text. If specified, an embodiment may depict the environment or setting. Such an environment or setting

20 might also be considered as a background upon or in which the text is portrayed. Consider, for example, the sentence: "*John walked through the forest.*" In this example, the environment or setting of the scene is a forest. The environmental database 38, of FIG. 1, may be used to supply

environments or settings.

However, there may be no explicit environment or setting. Consider, for example, the sentence: "*The flower is blue.*" Rather than depicting a blue flower floating on a page, it may be possible to supply an optional background setting. In the simplest case, this may be a ground plane and/or a supporting object. For more visually complex cases, it may be preferable to include more detail in the environment or scene. For example, an embodiment may put the flower in a vase on a fireplace mantle in the middle of a fully decorated living room.

### Abstract Relations

Scenes generated from input text are not limited to those scenes that may be directly depictable; abstract relations may also be depicted. FIG. 20 is one possible example of a rendered image 136 of the sentence: "*The cowboy is studying mycology.*" The idea of a "cowboy" may be represented by a male humanoid figure 138 wearing a cowboy style hat 140. The idea of "studying" may be represented by, for example, the humanoid figure 138 holding an open book 142. A book having a cover bearing the word "Mycology" 144 might, for example, represent the idea of the study of mycology. An image of a mushroom 146 may additionally be displayed on the book's 142 cover.

The concept of negation may also be represented. FIG. 21 is one possible example of a rendered image 148 of the sentence: "*The blue daisy is not in the vase.*" Negation may, for example, be represented by the superimposition of a circle with a diagonal line 150 over the image of a blue daisy 152 in a vase 154.

The above listing of abstract relations is not meant to be exclusive. Other abstract relations may be depicted without departing from the invention herein.

### Cardinality

5 The concept of cardinality may also be represented. FIG. 22 is one possible example of a rendered image 156 of the sentence: "*The two cats are on the table.*" Placing a first cat 158 next to a second cat 160 on a table 162 represents the concept of cardinality, the number of elements in a given set.

### Recursive Substitution

A scene may be generated and recursively placed into a second scene. As an example, consider the sentences "*John said that the cat was on the table. The animal was next to a bowl of apples.*" An embodiment of the system might construct a first scene of a cat on a table, the cat being next to a bowl of apples. The embodiment might construct a second scene of a male character having a speech bubble pointing toward his mouth. The first scene could be recursively inserted into the second scene to depict the sentences just considered. Other applications of recursive substitution are also available.

### Rendering Images

20 After a three-dimensional scene description has been generated, a three-dimensional image can be rendered using any number of three-dimensional rendering programs. A large number of rendering packages are available to do this. For example, graphics cards and game

consoles can do this in real time, techniques such as ray-tracing, which can incorporate shadows and diffraction, may also be used. Thus, a scene description generated by the system may be converted to a low-level scene description of, for example, the VRML type and then applied to a rendering package. Of course, the scene description generated by the system may alternatively 5 be converted to a low-level scene description of a type suitable for input to a video game. Other conversions to other types of low level scene descriptions are, of course, also available.

Though not properly part of scene creation, to render a scene, the scene needs to be illuminated.

This may be done by illuminating objects from different directions using virtual light sources.

### 10 **Viewing Perspective**

A user of the system may choose a camera position, or viewing perspective, from which to render the scene. Alternatively, heuristics may be added to automate the viewing perspective.

### 15 **Potential Uses**

A Text-to-Scene system may be integrated into any other system (e.g., any other software). In an embodiment, the system may be integrated into a three-dimensional video type game, such as a game on a Sony PlayStation 2®, or may be integrated into the firmware of the computer running the game. A user might enter text into the game in order to have the embedded system, for example, create a new room having user-described lighting and 20 environment or a new object having user-described attributes. The newly created images might become part of the game. Databases, such as those for three-dimensional models, might be accessed via a communication network, such as the Internet or an intranet. In another

embodiment, the system might be used as an aid in foreign language interpretation. A user might enter text into the system in the foreign language and view the scene that the user has described. If the user were to use the language incorrectly, or were to use unintended words, then the user would be able to determine that the image generated was not the image the user intended. In 5 another embodiment the system might reside on the Internet and be used to create, for example, greeting cards representative of the text entered by a user. Of course, the system may be used in many ways other than those listed here, without departing from the scope of the invention.

While there has been described what are believed to be the preferred embodiment of the present invention, those skilled in the art will recognize that other and further changes and modifications may be made thereto without departing from the spirit of the invention, and it is intended to claim all such changes and modifications as fall within the true scope of the invention.

10  
15  
20  
25  
30  
35  
40  
45  
50  
55  
60  
65  
70  
75  
80  
85  
90  
95